Introduction to Dynamic Programming

Prof. Dr. Aleksander Berentsen

Spring 2010

Adapted from the book "Dynamic Economics: Quantitative Methods and Applications" by Jerome Adda and Russell W. Cooper. The MIT Press (October 12, 2003).

1 Envelope theorem

The Envelope theorem describes how the maximal/minimal value of a variable changes, when the parameters of the model change.

Proposition 1 *Let*

$$v(a) = \max_x f(x, a)$$

Then:

$$\frac{dv(a)}{da} = \left.\frac{\partial f(x, a)}{\partial a}\right|_{x=x^*(a)}$$

where $x^*(a) = \arg\max_x f(x, a)$*. The envelope theorem tells us that we can ignore the partial effect of* x*.*

Proof.

$$v(a) \equiv f[x^*(a), a]$$

Deriving with respect to a on both sides gives:

$$\frac{dv(a)}{da} = \frac{\partial f[x^*(a), a]}{\partial x}\frac{\partial x^*(a)}{\partial a} + \frac{\partial f[x^*(a), a]}{\partial a}$$

Because $x^*(a)$ is the maximum of f, we know that:

$$\frac{\partial f[x^*(a), a]}{\partial x} = 0$$

Then:

$$\frac{dv(a)}{da} = \frac{\partial f[x^*(a), a]}{\partial a}$$

Which is the same as:

$$\frac{dv(a)}{da} = \left.\frac{\partial f(x, a)}{\partial a}\right|_{x=x^*(a)}$$

■

The Envelope theorem can be derived for the restricted optimization problem.

Proposition 2 *Let*

$$m(a) = \max_x f(x,a) \qquad \text{s.t.} \qquad g(x,a) = 0 \quad \text{and} \quad x \geq 0,$$

Let $\mathcal{L}(x,a,\lambda)$ be the corresponding Lagrange function, and let $x^*(a)$ and $\lambda^*(a)$ be the corresponding values, which solve the Kuhn-Tucker conditions. Then:

$$\frac{dm(a)}{da} = \left.\frac{\partial \mathcal{L}}{\partial a}\right|_{x^*(a),\lambda^*(a)}$$

2 The "Cake Eating" Example: Direct Solution

- Cake of the size W_1
- $t = 1, 2, ..., T$ periods
- In each period, a bit of the cake can be eaten.

How does the cake have to be eaten optimally across time? What is the optimal consumption?

c_t: Consumption in period t

$u(c_t)$: Utility in period t

Let $u(c_t)$ be derivable, strictly monotonic and stricly concave. Then $\lim_{c \to 0} u'(c) \to \infty$. This so-called Inada condition makes sure that we eat at least a little bit of the cake in each period (i.e., no corner solution is possible).

The houselhold's discounted utility is:

$$\sum_{t=1}^{T} \beta^{t-1} u(c_t)$$

The modification of the cake across time is:

$$W_{t+1} = W_t - c_t \qquad t = 1, 2, ..., T \tag{1}$$

Direct solution:

$$\max_{\{c_t\}_1^T, \{W_t\}_2^{T+1}} \sum_{t=1}^{T} \beta^{t-1} u(c_t) \tag{2}$$

$$\begin{aligned} \text{s.t.} \quad W_{t+1} &= W_t - c_t \qquad t = 1, 2, ..., T \\ c_t &\geq 0 \\ W_{T+1} &\geq 0 \end{aligned} \tag{3}$$

W_T is for the consumption at the last period, so W_{T+1} is the final stock.

The constraints (3) can be summarized as:

$$\sum_{t=1}^{T} c_t + \sum_{t=1}^{T} W_{t+1} = \sum_{t=1}^{T} W_t$$
$$\sum_{t=1}^{T} c_t + W_{T+1} = W_1 \tag{4}$$

Equation 4 means that the totality of consumption added to the final stock must be equal to the initial stock.

$$\max_{\{c_t\}_1^T, \{W_t\}_2^{T+1}} \sum_{t=1}^{T} \beta^{t-1} u(c_t) \qquad \text{s.t. (4), } c_t \geq 0 \text{ and } W_{T+1} \geq 0$$

Here, we do not take the constraint $c_t \geq 0$ into account because we assume that the marginal utility of consumption goes to infinity when the consumption of a period approaches zero. With such an inequality, we simply maximize and then check that the constraint is respected. In this case, it will be respected because of the Inada condition.

Maximization

$$\mathcal{L} = \sum_{t=1}^{T} \beta^{t-1} u(c_t) + \lambda[W_1 - \sum_{t=1}^{T} c_t - W_{T+1}] + \phi W_{T+1}$$

First-order conditions

1. For c_t:

   $$\beta^{t-1} u'(c_t) = \lambda \qquad t = 1, 2, ..., T$$

   λ: Lagrange multiplier of the constraint 4. This FOC allows us to equal two periods:

   $$u'(c_t) = \beta u'(c_{t+1}) \tag{5}$$

   which is the **Euler equation.**

2. For W_{T+1}:

   $$\lambda = \phi$$

   ϕ: Lagrange multiplier of the constraint $W_{T+1} \geq 0$. This FOC tells us that in order to respect the Kuhn-Tucker conditions $\phi W_{T+1} = 0$, the following expression must hold:

   $$W_{T+1} = 0$$

2.1 Interpretation of the Euler equation

At equilibrium, the marginal costs, when the consumption in one period is reduced, correspond exactly to the discounted marginal utility that is generated when the savings are consumed in the next period. When the Euler equation is satisfied, it is impossible to increase utility by a transfer of consumption in the next period.

As for the indirect utility function, $V_T(W_1)$ describes the maximal utility that a household can reach, when the cake has a size of W_1 and the number of periods is T. $V_T(W_1)$ denotes the **value function**:

$$V_T(W_1) = \max_{c_t} \sum_{t=1}^{T} \beta^{t-1} u(c_t) + \lambda[W_1 - \sum_{t=1}^{T} c_t - W_{T+1}] + \phi W_{T+1}$$

What happens to $V_T(W_1)$, when the size of the cake changes marginally?

$$\frac{dV(W_1)}{dW_1} = V_T'(W_1) = \lambda = \beta^{t-1} u'(c_t) \qquad t = 1, 2, ..., T$$

One can see that it does not depend on whether the additional cake is eaten, because the household is indifferent at equilibrium.

3 The "Cake Eating" Example: Dynamic Programming

3.1 Finite time horizon

We introduce the period 0:

$$V_{T+1}(W_0) = \max_{c_0} \ u(c_0) + \beta V_T(W_1) \tag{6}$$

with

$$W_1 = W_0 - c_0$$

where W_0 is the size of the cake at time 0.

We choose c_0 and thus directly define W_1. The function $V_T(W_1)$ gives us the expected discounted utility, when we have a cake of size W_1 at our disposal in period 1. We have already chosen this consumption optimally.

⇒ Principle of optimality (Richard Bellman)

First-order condition:

$$u'(c_0) = \beta V_T'(W_1)$$

where

$$V_T'(W_1) = u'(c_1) = \beta^t u'(c_{t+1}) \qquad t = 1, 2, ..., T-1$$

⇒ Euler equation

3.1.1 Example

Consider the following utility function $u(c) = \ln(c)$.

- Let $T = 1$, such that
$$V_1(W_1) = \ln(W_1)$$
- Let $T = 2$.
$$V_2(W_1) = \max_{W_2} \ \ln(W_1 - W_2) + \beta \ln(W_2)$$
First-order condition:
$$\frac{1}{c_1} = \frac{\beta}{c_2}$$
Restriction:
$$c_1 + c_2 = W_1$$
Therefore
$$c_1 = \frac{W_1}{1+\beta}\ , \quad c_2 = \frac{\beta W_1}{1+\beta}$$
Construction of the value function
$$V_2(W_1) = \ln\left(\frac{W_1}{1+\beta}\right) + \beta \ln\left(\frac{\beta W_1}{1+\beta}\right) \tag{7}$$

Remark 1*The* max*-Operator has been removed in (7), because we have already inserted the optimal values.*

$$V_2(W_1) = A_2 + B_2 \ln(W_1)$$
where
$$\begin{aligned} A_2 &= \ln\left(\frac{1}{1+\beta}\right) + \beta \ln\left(\frac{\beta}{1+\beta}\right) \\ B_2 &= (1+\beta) \end{aligned}$$

- Let $T = 3$.
(A)
$$V_3(W_1) = \max_{W_2} \ \ln(W_1 - W_2) + \beta V_2(W_2)$$
First-order condition:
$$\frac{1}{c_1} = \beta V_2'(W_2)$$
(B)
$$V_2(W_2) = \max_{W_3} \ \ln(W_2 - W_3) + \beta V_1(W_3)$$
First-order condition:
$$\frac{1}{c_2} = \beta V_1'(W_3)$$

Envelope condition:

$$V_2'(W_2) = \frac{1}{c_2}$$

(C)

$$V_1 = \ln(W_3) \qquad \text{because } W_4 = 0$$
$$\Rightarrow V_1'(W_3) = \frac{1}{c_3}$$

- Backward resolution:

  (C) and (B) imply
  $$\frac{1}{c_2} = \frac{\beta}{c_3}$$
  (B) and (A) imply
  $$\frac{1}{c_1} = \frac{\beta}{c_2}$$
  Resource restriction
  $$c_1 + c_2 + c_3 = W_1$$
  Therefore:
  $$c_1 = \frac{W_1}{1+\beta+\beta^2} \qquad c_2 = \frac{\beta W_1}{1+\beta+\beta^2} \qquad c_3 = \frac{\beta^2 W_1}{1+\beta+\beta^2}$$
  Construction of the value function
  $$V_3(W_1) = \ln c_1 + \beta\left[\ln c_2 + \beta \ln c_3\right].$$

3.2 Infinite time horizon

$$\max_{\{c_t\}_1^\infty, \{W_t\}_2^\infty} \sum_{t=1}^{\infty} \beta^{t-1} u(c_t)$$
$$W_{t+1} = W_t - c_t \qquad t = 1, 2, ...$$

Dynamic programming

$$V(W) = \max_{c \in [0,W]} u(c) + \beta V(W - c) \tag{8}$$

$V(W)$ expresses the maximal utility possible given the fact that I behave optimally from now on. $V(W-c)$ expresses the maximal utility possible given the fact that I behave optimally from tomorrow on.

Where W is defined as a state variable and c as a control variable.

Let

$$W' = W - c$$

We now can write:

$$V(W) = \max_{W' \in [0,W]} u(W - W') + \beta V(W') \tag{9}$$

$\Rightarrow$ **Functional equation** respectively **Bellman equation**.

To solve the problem, we have three steps. The idea is to look for a function $V(W)$ that fulfills equation 9 for all W.

Stationarity: The time index does not come up, because the solution $V(W)$ does not depend on the time.

1. First-order condition:
$$u'(c) = \beta V'(W')$$
How is $V'(W')$ calculated?

2. Envelope condition
$$V'(W) = u'(c) \Leftrightarrow V'(W') = u'(c')$$
c chosen optimally

3. FOC and EC put together:
$$u'(c) = \beta u'(c')$$
$\Rightarrow$Euler equation

Policy function
$$c = \phi(W) \quad \text{resp.} \quad W' = e(W) = W - \phi(W)$$
$$u'\left[\phi(W)\right] = \beta u' \left\{\phi\left[(W - \phi(W)\right]\right\}$$
$\Rightarrow$ important in empirical research.

3.2.1 Example

Normally, there is no way to find an explicit function for the value function $V(W)$. However, it is possible for $u(c) = \ln c$. We guess the solution and verify it in the end.
$$V(W) = A + B\ln(W)$$
This guess gives us
$$V(W) = \max_{W' \in [0,W]} \ \ln\left(W - W'\right) + \beta V(W')$$
$$\Rightarrow A + B\ln(W) = \max_{W' \in [0,W]} \ \ln(W - W') + \beta V(A + B\ln(W')) \tag{10}$$

First-order condition:
$$\frac{1}{W - W'} = \beta B \frac{1}{W'}$$
$$\Rightarrow \quad W' = \beta B W - \beta B W'$$
$$\Rightarrow \quad W' = \frac{\beta B W}{1 + \beta B}$$

Substitute in (10):

$$\begin{aligned} A + B\ln(W) &= \ln\left(\frac{W}{1+\beta B}\right) + \beta\left[A + B\ln\left(\frac{\beta B W}{1+\beta B}\right)\right] \\ &= \ln(W) - \ln(1+\beta B) + \beta A + \beta B\ln(W) + \beta B\ln\left(\frac{\beta B}{1+\beta B}\right) \\ &= \underbrace{(1+\beta B)}_{B}\ln(W) - \underbrace{[1+\beta B]\ln(1+\beta B) + \beta\left(A + B\ln(\beta B)\right)}_{A} \end{aligned}$$

$$\begin{aligned} 1+\beta B &= B \\ B &= \frac{1}{1-\beta} \end{aligned}$$

For this functional form, the first-order condition is:

$$\begin{aligned} \frac{1}{c} &= \beta\frac{1}{1-\beta}\frac{1}{W'} \\ &= \frac{\beta}{1-\beta}\frac{1+\frac{\beta}{1-\beta}}{\frac{\beta}{1-\beta}W} \\ &= \frac{1}{W(1-\beta)} \end{aligned}$$

$$\begin{aligned} &\Rightarrow \quad c = (1-\beta)W \\ &\Rightarrow \quad W' = c\frac{\beta}{1-\beta} = \beta W \end{aligned}$$

3.2.2 Taste shocks

Dynamic programing allows for uncertainty.

Example: Let $\varepsilon u(c)$ be the utility for a given period, where ε is a stochastic variable. The agent knows the shock when he chooses his consumption, but he does not know his future preferences.

Let $\varepsilon \in \{\varepsilon_h, \varepsilon_l\}$ with $\varepsilon_h > \varepsilon_l > 0$.

The shock follows a first-order Markov process. The probability that the shock takes a particular realization depends only on the realization in the last period.

π_{ij} is the probability that the value of ε skips in the current period from state i over to state j in the next period.

Transition matrix

		$t+1$: ε_h	$t+1$: ε_l
t	ε_h	π_{hh}	π_{hl}
	ε_l	π_{lh}	π_{ll}

$$\pi_{lh} \equiv \Pr(\varepsilon' = \varepsilon_h \mid \varepsilon = \varepsilon_l)$$

Cake eating:

$$V(W, \varepsilon) = \max_{W'} \; \varepsilon u(W - W') + \beta E_{\varepsilon'|\varepsilon} V(W', \varepsilon')$$

for all ε (i.e., $\varepsilon_h, \varepsilon_l$)
where $W' = W - c$

First-order condition

$$\varepsilon u'(W - W') = \beta E_{\varepsilon'|\varepsilon} V_1(W', \varepsilon') \qquad \varepsilon = \varepsilon_h, \varepsilon_l \tag{11}$$

Envelope condition

$$V_1(W, \varepsilon) = \varepsilon u'(W - W')$$

$$\Rightarrow V_1(W', \varepsilon') = E_{\varepsilon'|\varepsilon} \left[\varepsilon' u'(W' - W'')\right] \tag{12}$$

(11) and (12) imply

$$\varepsilon u'(W - W') = \beta E_{\varepsilon'|\varepsilon} \left[\varepsilon' u'(W' - W'')\right]$$

$\Rightarrow$ stochastic Euler equation

Policy function

$$W' = \varphi(W, \varepsilon)$$

Euler equation

$$\varepsilon u'(W - \varphi(W, \varepsilon)) = \beta E_{\varepsilon'|\varepsilon} \left[\varepsilon' u'(\varphi(W, \varepsilon) - \varphi(\varphi(W, \varepsilon), \varepsilon'))\right].$$